



Author: Marc Staimer, President Dragon Slayer Consulting

W H I T E P A P E R

Top 4 Reasons Oracle ZS7-2 is Superior to Isilon

Compelling Justification Why Isilon Users Should Switch

Top 4 Reasons Oracle ZS7-2 is Superior to Isilon

Oracle ZS7-2 versus DELL EMC Isilon Gen6 Introduction

Oracle ZS7-2 (ZS7-2) is a state-of-the-art, scale-up, multi-purpose file and block storage system. The ZS7-2 is the latest Oracle ZFS Storage Appliance. This new generation delivers considerably greater performance, improved ease-of-use, greater efficiencies, simplified management, and lower TCO. The ZS7-2 is exclusively co-engineered with the Oracle Database to deliver unmatched Oracle Database performance, reduced storage consumption, and greater data protection.

The DELL EMC Isilon Gen6 (Isilon) is a multi-node scale-out, multi-purpose file and block storage system designed for storing large amounts of unstructured data. Although it was not designed for high performance or transactional environments, DELL EMC have often sold it into them. The results have not been pretty with frequently unhappy users. The current generation made a concerted effort to noticeably improve performance. However, an examination of its results demonstrate it still cannot be considered highly performant or effective for transactional environments.

Both the ZS7-2 and Isilon Gen6 are vastly improved. A deeper dive makes it clear that the Oracle ZS7-2 has widened the gap over Isilon in performance and three other dimensions.

1. Performance
2. Flexibility
3. Oracle Database Co-Engineering
4. Much Lower TCO

Performance

DELL EMC Isilon (Isilon) has never been architected for performance. The Isilon engineering team have been trying to rectify that oversight since it first hit the market. Their results have been mixed. Despite constant efforts to improve performance via the storage media, nodal hardware, and system software they have not been able to overcome the limitations of OneFS – Isilon file system. OneFS is designed for large files and does not do very well with files smaller than 1 MB. It is a highly inefficient file system adding unnecessary latency to every IO. Hardware and storage stack improvements can only mask an inadequate amount of the OneFS inefficiencies. They cannot make it go away. Another frequently reported user performance issue is storage pool overflow. A higher performance pool fills up, it can overflow into a lower performance pool, automatically causing user frustration.



Oracle ZS7-2 Storage (ZS7-2) is architected from the ground up for exceptional IO and throughput performance with extremely low latency. It starts with the exceedingly efficient and production-hardened ZFS (ZS7-2 file system). ZFS is the award-winning file system utilized by dozens of storage and systems vendors worldwide; the Oracle version contains many more security and enterprise features. Next is the DRAM-centric architectural design and symmetric multi-processing (SMP) operating system, which makes highly effective utilization of up to 3 TB of intelligent DRAM for caching and nearly a hundred multi-core CPUs. DRAM is up to 1000x faster and lower latency than flash SSDs and over 100x faster than NVMe SSDs. ZS7-2 DRAM caching accelerates IO and throughput. The ZS7-2 architecture is an **"AND"** not an **"OR"** architecture. It enables high IOPS **AND** high throughput. It pools effectively up to 3.7 PB of flash SSDs **AND** 11.5 PB of HDD. This makes the ZS7-2 ideal for those applications that need the **"AND"**. Applications such as: Oracle Database and application workloads, private and public cloud deployments, visual effects (VFX) digital rendering, electronic design automation (EDA), video streaming, software development, target backup and recovery storage, and even virtual machine (VM) storage.

Isilon is an **OR** architecture. For example: a node/chassis can be all-flash **OR** HDD. While different drive types are supported, they cannot be mixed in the same node or chassis. It is one type **OR** another of the drives. There are six unique node/chassis types. Each has different drives and memory (DRAM caching).

How the ZS7-2 compares in performance to Isilon depends on several factors including the Isilon node/chassis type, drives utilized, and how the Isilon published performance numbers are interpreted. For the sake of argument, the Isilon F800 all-flash chassis is used for comparison.

Although DELL EMC publishes aggregate chassis and a maxed out 36 chassis system streaming throughput performance for the all-flash F800, they do not publish IOPS. Aggregate falsely implies jobs can stripe across all 4 nodes within a chassis or even across all 36 chassis in a system. They don't. The aggregate or concurrent throughput performance is the total topline Isilon performance available for multiple jobs or streams. But any specific job or stream is limited to that of a single node. Isilon's F800 configured topline potential read throughput is 3.75 GB/s per node; 15 GB/s per chassis; and 540 GB/s per maxed out F800 system. The write throughput topline is 1.9 GB/s per node; 7.6 GB/s per chassis; and 273.6 GB/s per maxed out F800 system.

The ZS7-2 **AND** architecture enables the active-active controllers to provide the aggregate performance to a specific job such as RMAN. RMAN recovery performance (reads) from a ZS7-2 is 66.3 TB/hr or 18.42 GB/s (9.208 GB/s per node). Just comparing the node performance equates into more than 2.4x the max Isilon read throughput. RMAN backup performance (writes) from the ZS7-2 is 54.8 TB/hr or 15.22 GB/s (7.611 GB/s per node). That's more than 4x the max Isilon write throughput per node. What this means is the ZS7-2 will provide a lot more usable throughput performance to any given job or stream while the Isilon system – not the chassis – has more potential total performance available.

The benchmarks have been proving out the ZS7-2 exceptional performance. Preliminary¹ SpecSFS2014 SP2 Database results (this workload represents the typical behavior of a transactional SQL database²), SpecSFS2014 SP2 VDA results with Spectre/Meltdown patch (the business metric for the VDA benchmark is STREAMS³), SpecSFS2014 SP2 SWBUILD results (software build type workload is a classic meta-data intensive build workload⁴), SpecSFS2014 SP2 EDA results (the EDA workload is representative of the Electronic Design Automation environments⁵), all demonstrate unmistakable ZS7-2 best-in-class performance.

DELL EMC does not publish any Isilon industry performance benchmark results. No matter DELL EMC assertions or propaganda, the ZS7-2 delivers superior performance results over Isilon.

Flexibility

The ZS7-2 is designed with flexibility and ease-of-use in mind. The DRAM, SSDs, HDDs, network connections, HBAs, are all flexible and configurable based on requirements. SSDs and HDDs can be mixed in the same system. Different drive types and capacities can be mixed in the same system. Infiniband (40 Gbps) and Ethernet (10/40 Gbps) can be mixed in the same system. The number of network connections are variable and configurable. Multiple different types of workloads can be mixed within the same system.

Isilon is not as flexible. As previously mentioned, drive types and different capacities cannot be mixed within a node or chassis. They can only be mixed within a system with multiple different chassis. But it tends to create operational and management complexity while complicating performance management. It is doable and supported; however, best practices recommend different chassis types be deployed in their own clustered systems. Data can be moved or migrated from the higher performing systems to the lower performing ones based on tiering and replication policies. Although this simplifies individual system management, it significantly complicates ecosystem management and increases cost.

¹ Preliminary means the performance numbers have not yet been certified for publication.

² The complete workload is a mixture of DB_TABLE and DB_LOG workloads. The DB_TABLE workload is the database component, and DB_LOG represents the log writer component of a database operation.

³ The benchmark consists of two workload objects: VDA1 (data stream) and VDA2 (companion applications). Each stream corresponds to a roughly 36 Mb/s bit rate, which is in the upper range of high definition video. The goal of the storage admin is to provide as many simultaneous streams as possible while meeting the bit rate and fidelity constraints.

⁴ This workload was derived from analysis of software builds, and traces collected on systems in the software build arena. Conceptually, these tests are similar to running UNIX 'make' against several tens of thousands of files. The file attributes are checked (metadata operations) and if necessary, the file is read, compiled, then data is written back out to storage.

⁵ This workload is based on network traces collected from real environments and input from domain experts in and published documentation of the real-world implementation of the EDA applications being simulated. This workload also is the first workload in the SPEC SFS 2014 suite to have data sets that are both compressible and dedupable. The addition of support for dedupable data sets is a new feature within the SPEC SFS 2014 SP2 benchmark release.

One more thing that frustrates many Isilon users is that some upgrades require the entire system cluster to be rebooted simultaneously. That should simply not be necessary in this day and age.

[Oracle Database Co-Engineering](#)

The ZS7-2 is co-engineered with the Oracle Database. The results of that co-engineering are 3 exclusive ZS7-2 capabilities of Oracle Intelligent Storage Protocol version 2 (OISP2), ZS7-2 Hybrid Columnar Compression (HCC) integration, and the tight correlation of ZS7-2 storage analytics with the Oracle Database Automated Workload Repository (AWR).

OISP2 is unique to the ZS line of Oracle storage. It is a direct communication pipeline between the Oracle Database and the ZS7-2. That pipeline has a profound impact on Oracle Database performance, and storage management that makes the job of the DBA infinitely easier.

- For low latency IO-sensitive operations, OISP2 prioritizes those IOs first, cuts latency on critical IOs by up to 13.5x, and runs up to 19% faster for transaction-intensive applications.
- For higher throughput cache bursting, Oracle Database operations such as backup and restore, it runs up to 33% faster, while creating optimal cache efficiency for other applications.
- For storage management, OISP2 auto-sets up and auto tunes the ZS7-2, reducing steps and processes by as much as 70%.

The integration of Oracle Database HCC with ZS7-2 produces 10-50x compression and an average of 12x. That compression greatly increases data warehouse and archive efficiency while radically reducing Oracle Database storage capacity consumption. It also speeds up Oracle Database processing of the compressed data because unlike Isilon deduplication, it does not require rehydration to be read. The average data reduction alone for the ZS7-2 is 3-6x better⁶ than the best Isilon Oracle Database storage reduction results.

Oracle Database AWR integration with the ZS7-2 storage analytics makes the DBA life much simpler when it comes to performance troubleshooting. The tight correlation provides the DBA unmatched actionable insights into IO storage performance on an IO type basis enabling them unprecedented control.

Isilon has nothing like OISP2, HCC or any Oracle Database integration other than RMAN compatibility and never will. The question for any Oracle Database administrator becomes quite simple, do they want storage that accelerates their database, reduces their workload, reduces their storage consumption, and reduces their costs from reduced storage requirements? Or do they want to get poorer Oracle Database results, work harder, and spend more budget? This should be a rhetorical question.

[Much Lower TCO](#)

The ZS7-2 has several ways in which it reduces costs over Isilon. As just discussed, the Oracle Database consumes on average 75 - 83% less storage capacity on the ZS7-2 than on Isilon because of HCC. Less storage capacity requirements equals reduced capital costs and reduced operating costs (maintenance, power, cooling, rack space, and more)

New improvements in RAIDZ (2+1) enables the ZS7-2 to run very close to mirroring (2+2) performance with a 25% reduction in drives. Isilon cannot do that. Fewer parity drives equals more efficiency and lower cost. Isilon is known to be highly inefficient for many small files such as millions of files less than 128 KB in size such as in a CAD/CAM. Small files cause the Isilon protection level to utilize mirroring which comparatively to RAIDZ, requires 33% more capacity and cost.

ZS7-2 Oracle Database co-engineering means reduced admin time spent on storage configuration, tuning, and troubleshooting. Reduced time equals reduced cost and higher admin productivity. Once again, Isilon lacks this co-engineering.

Noticeably better performance means faster job completions, increased productivity, and faster time-to-market. Faster-time-to-market means decreased costs and/or increased revenues.

⁶ Users report that Isilon deduplication is considerably less effective than other platforms and not very effective at all on Oracle Databases getting at most a 2.0 to 2.8X data reduction.

All of these cost savings and revenue gains come before the system is priced out and before any discounts. The ZS7-2 definitively delivers a much lower TCO than Isilon.

Conclusion

Both the Oracle ZS7-2 and DELL EMC Isilon Gen6 have made significant improvements in their latest generations of storage. But, the ZS7-2 has improved considerably more than Isilon. Better performance, more flexibility, unique co-engineered integration with the Enterprise database leader – Oracle Database, and significantly lower TCO makes the ZS7-2 a compelling “no-brainer” choice over Isilon.

[For More Information on the Oracle ZS7-2](#)

Go to: <https://www.oracle.com/storage/nas/zs7-2>

Research sponsored by Oracle. **About Dragon Slayer Consulting:** Marc Staimer, as President and CDS of the 20-year-old Dragon Slayer Consulting in Beaverton, OR, is well known for his in-depth and keen understanding of user problems, especially with storage, networking, applications, cloud services, data protection, and virtualization. Marc has published thousands of technology articles and tips from the user perspective for internationally renowned online trades including many of TechTarget's Searchxxx.com websites and Network Computing and GigaOM. Marc has additionally delivered hundreds of white papers, webinars, and seminars to many well-known industry giants such as: Brocade, Cisco, DELL, EMC, Emulex (Avago), HDS, HPE, LSI (Avago), Mellanox, NEC, NetApp, Oracle, QLogic, SanDisk, and Western Digital. He has additionally provided similar services to smaller, less well-known vendors/startups including: Asigra, Cloudtenna, Clustrix, Conduvis, DH2i, Diablo, FalconStor, Gridstore, Nexenta, Neuxpower, NetEx, NoviFlow, Pavilion Data, Permabit, Qumulo, SBDS, StorONE, Tegile, and many more. His speaking engagements are always well attended, often standing room only, because of the pragmatic, immediately useful information provided. Marc can be reached at marcstaimer@me.com, (503)-312-2167, in Beaverton OR, 97007.